



Europäisches Patentamt
European Patent Office
Office européen des brevets



Publication number: **0 592 117 A2**

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: **93307297.7**

Reg.

(51) Int. Cl.⁵: **G06F 9/46**

(22) Date of filing: **16.09.93**

(30) Priority: **24.09.92 US 951488**

(43) Date of publication of application:
13.04.94 Bulletin 94/15

(84) Designated Contracting States:
DE ES FR GB IT

(71) Applicant: **AMERICAN TELEPHONE AND
TELEGRAPH COMPANY**
32 Avenue of the Americas
New York, NY 10013-2412 (US)


(72) Inventor: **Allen, James Riley**
1364 Deep Run Road
Naperville, Illinois 60540 (US)

(74) Representative: **Buckley, Christopher Simon
Thirsk et al**
AT & T (UK) LTD., AT & T Intellectual Property
Division, 5 Mornington Road
Woodford Green, Essex IG8 0TU (GB)

(54) Asynchronous inter-process communications arrangement

(57) To overcome the lack of support in time-sharing and uniprocessor operating systems such as the UNIX® operating system for real-time and multiprocessor applications, there is provided an asynchronous inter-process communications capability that can be grafted onto the operating systems. Communicating processes (100, 103) communicate via datagram messages through logical asynchronous inter-process communications links (110) each having a synchronous segment (101) and an asynchronous segment (102). The links include a message-serving hub process (11) that communicates in a synchronous, buffer and semaphore based, manner with processes (100) that are message senders, and communicates in an asynchronous, queue and signals based, manner with processes (103) that are message destinations. The hub process may be implemented at any process level of the operating system.

FIG. 2



Technical Field

This invention relates generally to computer operating systems, and relates specifically to inter-process communication arrangements of operating systems.

Background of the Invention

Inter-process communication arrangements of computer operating systems come in two basic varieties: synchronous and asynchronous. Synchronous inter-process communications are those that require a common occurrence, such as common timing, a predefined relationship, or in-step operation, in order for two or more processes to communicate. Conversely, asynchronous inter-process communications are those that occur without a regular or predictable time relationship or correlation between the communicating processes. Another way of viewing the difference between synchronous and asynchronous inter-process communications is that synchronous communications require the transmitting process to coordinate its transmissions in time with reception by the receiving process or processes, whereas asynchronous communications allow the transmitting and receiving processes to transmit and receive independently in time of each other's operations.

On one hand, operating systems designed to support uniprocessor time-sharing operations, such as the UNIX® operating system, generally implement synchronous inter-process communications. On the other hand, multiprocessor and real-time operations typically require asynchronous inter-process communications for their support. Synchronous communications arrangements are generally difficult to use with a multiplicity of processes and too prone to deadlock to provide effective support for multiprocessor and real-time operations.

Unfortunately, this requirement undesirably limits the uses for popular operating systems, such as the UNIX system, that do not provide effective asynchronous inter-process communications capabilities. The advantages of such popular operating systems include a large supply of designers experienced in their use, who may be readily drawn upon to staff new projects and to get these projects substantively under way beyond the initial start-up, learning, curve. The advantages further include a large body of existing programs that may be drawn upon as a library and that may be re-used in new projects without having to be designed anew.

Of course, conventional operating systems often may be redesigned to support real-time or multiprocessor operations and to include asynchronous inter-process communications mechanisms. But conventionally, such redesign effectively results in a new operating system that destroys the above-enumerated

advantages of the system on which it is based. An illustrative example of such a redesigned operating system is the UNIX-RTR operating system, which is a UNIX-based system that supports real-time applications such as telephony switching.

Consequently, what the art requires is an asynchronous inter-process communications mechanism that can be grafted onto an operating system, such as the UNIX system, without adversely affecting any of its conventional interfaces and operational characteristics. However, the standard operating system facilities provided by the UNIX operating system and its clones and derivatives make this difficult to achieve. The operating systems' inter-process-related system calls (e.g., shared memory, semaphores, inter-process messages and signals) are synchronous in nature and do not lend themselves to the support of asynchronous communication schemes. In addition the UNIX-type operating systems do not normally provide a facility for inhibiting process switches. This creates the possibility of multiple processes accessing shared memory simultaneously, thereby greatly limiting the usefulness of shared memory.

Summary of the Invention

This invention is directed to solving the problems and meeting the needs of the prior art. According to the invention, there is provided an asynchronous inter-process communications arrangement that may be grafted onto an operating system (for example, onto the UNIX operating system) in order to overcome the limited support for real-time and multiprocessor applications. The asynchronous inter-process communications arrangement comprises a synchronous communications path segment for communicating with a first process, such as the message-sending process, of a pair of communicating processes within a single processor, an asynchronous communications path segment for communicating with a second process of the pair, and an arrangement, such as a procedure, for interconnecting the synchronous and the asynchronous communications path segments to form a communications path extending between the processes of the pair.

Illustratively, the synchronous segment is semaphore-based. It utilizes a buffer implemented in shared memory and shared by the plurality of message-sending processes. Access to the buffer by the sending processes and the interconnecting procedure is serialized by a semaphore having three states or values. Also illustratively, the asynchronous segment is buffer-based. It utilizes circular buffers implemented in shared memory, a different one for each message-receiving process. A write pointer to each buffer is controlled by the interconnecting procedure, while a read pointer to each buffer is controlled by the corresponding message receiving process. The pro-

cedure notifies message-receiving processes of presence of a message via a signals facility.

As required by the art, the invention provides an asynchronous inter-process communications mechanism that can be grafted onto an operating system without adversely affecting any of the operating system's conventional interfaces and operational characteristics. It thus provides support for both real-time and multiprocessor applications without destroying existing characteristics, features or advantages of the existing operating system. It is simple to implement; illustratively, it is implementable as just another process. Furthermore, it may be incorporated into any process level of the operating system.

These and other advantages and features of the invention will become apparent from the following description of an illustrative embodiment of the invention taken together with the drawing.

Brief Description of the Drawing

FIG. 1 is a logical diagram of a computer that incorporates an illustrative embodiment of the invention;

FIG. 2 is a block diagram of an inter-process communications link of the computer of FIG. 1;

FIG. 3 is a block diagram of synchronous data-structures-and-functions block of a hub function of the link of FIG. 2;

FIG. 4 is a flow diagram of a synchronous procedure of a sending process of the link of FIG. 2;

FIG. 5 is a flow diagram of a synchronous portion of a procedure of the hub function of the link of FIG. 2;

FIG. 6 is a block diagram of asynchronous data-structures-and-functions block of the hub function of the link of FIG. 2;

FIG. 7 is a flow diagram of an asynchronous portion of the procedure of the hub function of the link of FIG. 2; and

FIG. 8 is a flow diagram of an asynchronous procedure of a receiving process of the link of FIG. 2.

Detailed Description

FIG. 1 is a logical diagram of a computer 10. Computer 10 comprises one or more partitions 1-N. Computer 10 may be a multiprocessor system wherein each partition 1-N represents a separate processor. Or, computer 10 may be a uniprocessor system wherein partitions 1-N are merely logical groupings of functions or processes of the system. A uniprocessor system 10 need not be partitioned, i.e., it may consist of only a single partition.

Each partition 1-N includes a plurality of conventional processes 0-K, where the value of K may differ from partition to partition. Processes 0-K in a single

partition 1-N may communicate with each other and with processes in other partitions 1-N through conventional inter-process communication arrangements provided by the operating system of their individual partition 1-N. Therefore, all existing functionality is preserved. Illustratively, either each partition 1-N executes its own copy of the UNIX operating system, or all partitions 1-N execute a common copy of the operating system, as is conventional.

According to the invention, each partition 1-N further includes a hub function 11 that provides asynchronous inter-process communications capability among processes 0-K of that partition 1-N, and provides such capability between processes 0-K of that partition and the processes 0-K of the other partitions 1-N through associated one or more interface nodes 12. The interface nodes 12 are directly interconnected by communication paths 15. Illustratively, each path 15 is a conventional communication mechanism, such as a TCP/IP-protocol link on an Ethernet® LAN, or a serial point-to-point modem connection. Illustratively, each interface node 12 is implemented as a process in the corresponding partition 1-N. Implementation of interface nodes 12 as processes enables hub functions 11 to treat inter-partition communications, that is, communications to and from interface nodes 12, identically to intra-partition communications between processes 0-K.

In each partition 1-N, hub function 11 is also implemented as a process, at any one of the process levels provided by the operating system of that partition 1-N. For example, in the UNIX operating system, hub function 11 may be implemented either as a user-level (application) process, or as a kernel-level process, internal to the operating system. Intermediate or interpreter-level (shell) process implementations may also be envisioned.

As shown in FIG. 2, hub function 11 implements a logical asynchronous inter-process communications path or link 110 between each communication-sending process 100 and a communication-receiving process 103. Processes 100 and 103 are from the set of processes 0-K and interface nodes 12. Each inter-process communications link 110 is made up of two portions or segments: a synchronous inter-process communications segment 101 between sending process 100 and hub function 11, and an asynchronous inter-process communications segment 102 between hub function 11 and receiving process 103. A single synchronous segment 101 may form a part of a plurality of inter-process communications links 110, such as when sending process 100 is broadcasting to a plurality of receiving processes 103.

Hub function 11 acts as a message server between sending processes 100 and receiving processes 103. Hub function 11 comprises two sets of data structures and facilities 200 and 202, which are shown in FIGS. 3 and 6, respectively, and a procedure

201 which is diagramed in FIGS. 5 and 7. Synchronous segment 101 is implemented by a sending process 100 procedure which is diagramed in FIG. 4, by data structures and facilities 200, and the portion of procedure 201 which is diagramed in FIG. 5. Asynchronous segment 102 is implemented by the portion of procedure 201 which is diagramed in FIG. 7, by data structures and facilities 202, and by a receiving process 103 procedure which is diagramed in FIG. 8. Procedure 201 therefore acts as an interface between segments 101 and 102 that joins the two segments together to form link 110.

As shown in FIG. 3, data structures and facilities 200 comprise a message buffer 302 implemented in shared memory 301, and an associated semaphore 300. Shared memory 301 is a mechanism supported by the standard UNIX operating system which allows the plurality of processes 0-K, interface nodes 12, and hub function 11 to map a part of their own address spaces into the same portion 302 of physical memory. Semaphore facility 304 is also a mechanism supported by the standard UNIX operating system which allows a read-modify-write (RMW) sequence of operations to be performed as an atomic operation on a memory location. Semaphore facility 304 includes semaphore 300 and an associated semaphore queue 303. Semaphore 300 takes on three logical states or values: "0", "1", and "3". Semaphore 300 is initialized with a value of "3". As is conventional in the UNIX operating system, any process that attempts to perform an RMW sequence on semaphore 300 that would change the semaphore's value to a negative value is automatically put to sleep by the UNIX operating system on semaphore queue 303, and is awakened by the operating system only when it both (a) reaches the head of queue 303 and (b) execution of its previously-attempted RMW sequence will not change the semaphore's value to a negative value. Semaphore 300 implements what is conventionally referred to as a "critical region" around message buffer 302; it ensures that only one process has access to message buffer 302 at any one time.

When a process 0-K or an interface node 12 wishes to communicate with another process 0-K or an interface node 12 by means of an inter-process communications link 110, it executes the procedure of FIG. 4 to become sending process 100. Upon its invocation, at step 400, the procedure attempts to perform an RMW sequence that decrements the value of semaphore 300 by two, at step 401. The only time that the attempt at the RMW sequence succeeds is when semaphore 300 has a value of "3", which signifies that message buffer 302 is empty and available for use. If the value of semaphore 300 is "1" or "0", the RMW sequence blocks or suspends because the sequence would reduce the value of semaphore 300 to a negative value, and sending process 100 is put to sleep on queue 303, at step 402. When this sending process

100 reaches the head of queue 303 and semaphore 300 re-acquires a value of "3", sending process 100 is awakened, at step 403, and it returns to step 401 to attempt again the previously-failed RMW sequence.

If and when the RMW sequence of step 401 succeeds, the value of semaphore 300 becomes "1", which signifies that message buffer 302 is in use and being filled. Following its success at step 401, the procedure of FIG. 4 transfers into message buffer 302 whatever message sending process 100 wishes to send, at step 404. Illustratively, the message has the form of a datagram. The message includes an identifier of one or more receiving processes 103 which are intended to receive the message. The procedure then attempts to perform an RMW sequence that decrements the value of semaphore 300 by one, at step 405. Since steps 401 and 405 are always performed in the same order, this attempt never fails except in case of an error. Therefore, the value of semaphore 300 becomes "0" at step 405, which signifies that message buffer 302 is in use and has been filled with a message. Its task is completed, and the procedure of FIG. 4 now returns to the point of its invocation, at step 406.

Procedure 201 of hub function 11 starts executing upon initialization of its corresponding partition 1-N. Upon being started, at step 500 of FIG. 5, procedure 201 reads semaphore 300, at step 501, to determine if its value is "0", at step 502. If not, it means that a message is not stored in message buffer 302, and so procedure 201 merely returns to step 501 to await arrival of a message. Alternatively, to avoid consuming CPU resources, procedure 201 goes to sleep on semaphore 300, in a conventional manner, and is awakened when the value of semaphore 300 becomes "0". If the value of semaphore 300 is "0", it means that a message is stored in message buffer 302, and so procedure 201 retrieves the stored message, at step 503. This frees message buffer 302 for use by another sending process 100. Procedure 201 so indicates by performing an RMW sequence on semaphore 300 that increments the value of semaphore 300 by three, i.e., back to a value of "3". Procedure 201 then advances to step 700 of FIG. 7.

Turning to FIG. 6, data structures and facilities 202 comprise a plurality of circular data queues 602 implemented in shared memory 301, one queue 602 for each process 0-K and interface node 12, i.e., one for each entity that may be a receiving process 103.

The implementation of queues 602 must be sensitive to the fact that the UNIX operating system does not support any mechanism for preventing process switches. Consequently, switches between processes 201 and 103 may occur at any time, including when the running process is in the midst of operating on a queue 602. One way of dealing with this situation is to permit only two reference structures to a queue

602—a read index and a write index—and allow each reference structure to be operated on exclusively by only one of the processes 201 and 103. Also, it must be ensured that all read and write operations on the read and write indexes are atomic operations. For example, on some machines, the read and write indexes must each be limited to a single byte in length. Accordingly, each queue 602 comprises a plurality of entries 605, and has an associated write (W) pointer 603 under control of procedure 201 of hub function 11 and a read (R) pointer 604 under control of the corresponding receiving process 103. Procedure 201 writes messages into queues 602. Receiving processes 103 retrieve messages from queues 602, each from its own corresponding queue 602. Also included among data structures and facilities 202 is the conventional UNIX system signals facility 600. It is a mechanism by which the operating system notifies processes of the occurrence of asynchronous events, e.g., traps and interrupts.

After it retrieves a message from message buffer 302 in FIG. 5, procedure 201 determines from the message contents the identity of receiving process 103 to which the message is addressed, at step 700. Procedure 201 then checks the receiving process' corresponding queue 602 to determine if it is full, at step 701, in a conventional manner. If so, procedure 201 discards the message, at step 703, and then returns to step 501 of FIG. 5. It is not necessary to inform sending process 100 of the failure to deliver the message, because a datagram protocol is being used which does not guarantee message delivery. Under a datagram protocol convention, message retransmission is the responsibility of the application.

If procedure 201 determines at step 701 that queue 602 of receiving process 103 is not full, it writes the retrieved message into an entry 605 of that queue 602 which is pointed to by that queue's associated W pointer 603, at step 705, and increments W pointer 603, at step 706. Procedure 201 then checks whether this is the only message in that queue 602, i.e., whether that queue 602 had been empty, at step 707, again in a conventional manner. If so, procedure 201 causes signals facility 600 to send a signal to receiving process 103 to inform it of presence of a message in its queue 602, at step 708. Procedure 201 then returns to step 501 of FIG. 5.

In response to receiving a signal that was caused by procedure 201 to be sent at step 708, receiving process 103 eventually executes the procedure of FIG. 8. Upon its invocation, at step 800, the procedure accesses the process' corresponding queue 602 and retrieves a message from entry 605 pointed to by that queue's R pointer 604, at step 801. The procedure also increments R pointer 604, at step 802. The procedure then checks whether that queue 602 is empty, at step 803, in a conventional manner. If not, the procedure returns to step 801 to retrieve another mes-

sage from that queue 602. But if queue 602 is now empty, the procedure clears the signal that caused its invocation, at step 804, and then returns to the point of its invocation, at step 805. The inter-process communication between sending and receiving processes is thus completed.

Of course, it should be understood that various changes and modifications to the illustrative embodiment described above will be apparent to those skilled in the art. For example, a wide range of inter-processor communications schemes may be used—the described technique is independent of the inter-processor communications mechanism. Also, the multi-processor may be implemented as a single machine within a common cabinet, or it may be a network of discrete machines. Furthermore, different queuing data structures other than circular buffers may be used. Further yet, a UNIX inter-process message facility may be substituted for semaphores or shared memory. Such changes and modifications can be made without departing from the spirit and the scope of the invention and without diminishing its attendant advantages. It is therefore intended that all such changes and modifications be covered by the following claims.

Claims

1. An inter-process communications arrangement comprising:
 - a synchronous communications path segment for communicating with a first process of a pair of communicating processes within a single processor,
 - an asynchronous communications path segment for communicating with a second process of the pair, and
 - means for interconnecting the synchronous and the asynchronous communications path segments to form a communications path extending between the processes of the pair.
2. The arrangement of claim 1 wherein the interconnecting means comprise:
 - first means for communicating with the first process in a synchronous manner across the synchronous communications path segment to receive from the first process a communication destined for the second process; and
 - second means for communicating with the second process in an asynchronous manner across the asynchronous communications path segment to transfer the received communication to the second process.
3. The arrangement of claim 1 wherein:
 - the synchronous communications path

segment comprises

a shared-memory message buffer, and
 a semaphore means for creating a critical
 region around the message buffer; and
 the asynchronous communications path
 segment comprises
 a shared-memory message queue having
 a read pointer and a write pointer, and
 a signals means for giving notice of pres-
 ence of a message in the queue.

4. The arrangement of claim 1 CHARACTERIZED
 IN THAT

the arrangement is implemented using
 UNIX system facilities in a UNIX system operat-
 ing environment.

5. The arrangement of claim 1 wherein:

the synchronous communications path in-
 cludes

a message buffer for use by a plurality of
 message-sending processes,

a semaphore associated with the mes-
 sage buffer to control access to the message buf-
 fer and having three possible states, and

first operative means associated with the
 first process for performing a first operation on
 the semaphore when the semaphore is in a first
 state indicating that the buffer is available for use
 by the message-sending processes to place the
 semaphore in a second state thereby reserving
 the message buffer for exclusive use by the first
 process, and responsive to the reservation for
 storing a message in the message buffer and
 performing a second operation on the sema-
 phore to place the semaphore in a third state
 thereby freeing the message buffer for use by the
 interconnecting means; and

the interconnecting means include

second operative means responsive to the
 semaphore being in the third state for retrieving
 the stored message from the message buffer and
 performing a third operation on the semaphore to
 place the semaphore in the first state.

6. The arrangement of claim 1 wherein:

the asynchronous communications path
 includes

a plurality of message buffers each asso-
 ciated with a different message-receiving proc-
 ess;

means for sending a signal to a selected
 message-receiving process, and

first operative means associated with the
 second process and responsive to receipt of a
 signal from the signal-sending means for retriev-
 ing a stored message from the one of the mes-
 sage buffers associated with the second proc-

ess; and

the interconnecting means include

second operative means for storing a re-
 ceived message in the one of the message buf-
 fers associated with the second process and
 causing the signal-sending means to send a sig-
 nal to the second process.

7. The arrangement of claim 1 wherein:

the synchronous communications path in-
 cludes

a message buffer for use by a plurality of
 message-sending processes,

a semaphore associated with the mes-
 sage buffer to control access to the message buf-
 fer and having three possible states, and

first operative means associated with the
 first process for performing a first operation on
 the semaphore when the semaphore is in a first
 state indicating that the buffer is available for use
 by the message-sending processes to place the
 semaphore in a second state thereby reserving
 the message buffer for exclusive use by the first
 process, and responsive to the reservation for
 storing a message in the message buffer and
 performing a second operation on the sema-
 phore to place the semaphore in a third state
 thereby freeing the message buffer for use by the
 interconnecting means;

the interconnecting means include

second operative means responsive to the
 semaphore being in the third state for retrieving
 the stored message from the message buffer and
 performing a third operation on the semaphore to
 place the semaphore in the first state;

the asynchronous communications path
 includes

a plurality of message buffers each asso-
 ciated with a different message-receiving proc-
 ess;

means for sending a signal to a selected
 message-receiving process, and

third operative means associated with the
 second process and responsive to receipt of a
 signal from the signal-sending means for retriev-
 ing a stored message from the one of the mes-
 sage buffers associated with the second proc-
 ess; and

the interconnecting means further include

fourth operative means responsive to re-
 trieval of a message by the second operative
 means for storing the retrieved message in the
 one of the message buffers associated with the
 second process and causing the signal-sending
 means to send a signal to the second process.

8. An inter-process communications method com-
 prising the steps of:

synchronously communicating with a first process of a pair of communicating processes within a single processor to receive a communication from the first process for the second process; and

5

in response to receipt of the communication, asynchronously communicating with a second process of the pair to transfer the received communication to the second process.

10

9. The method of claim 8 wherein

the step of synchronously communicating comprises the step of

communicating with the first process through a shared-memory message buffer by using a semaphore facility to create a critical region around the message buffer; and

15

the step of asynchronously communicating comprises the step of

communicating with the second process through a shared-memory message queue, having a read pointer and a write pointer, by using a signals facility to indicate to the second process presence of a message in the queue.

20

25

10. The method of claim 8 CHARACTERIZED IN THAT

the method is implemented using UNIX system facilities in a UNIX system operating environment.

30

35

40

45

50

55

FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5

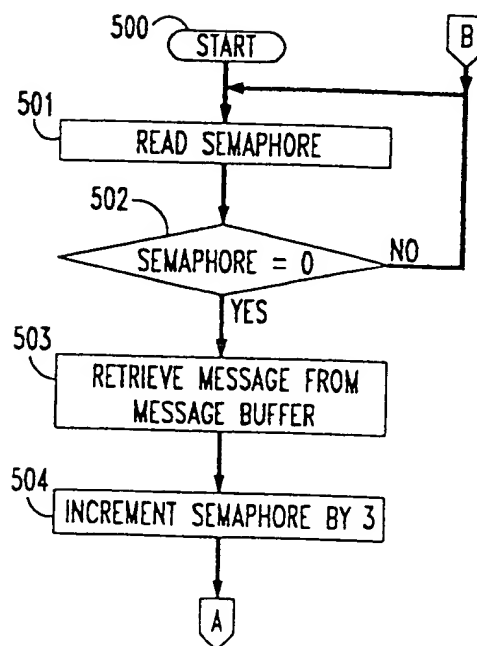


FIG. 6




FIG. 7

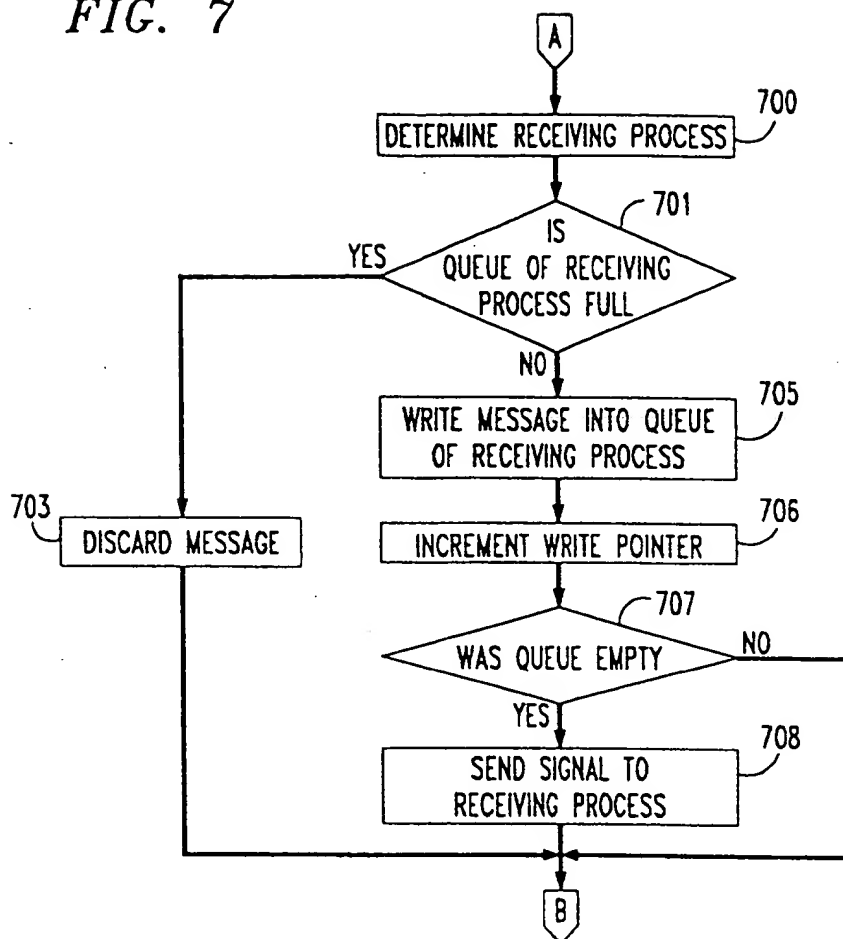
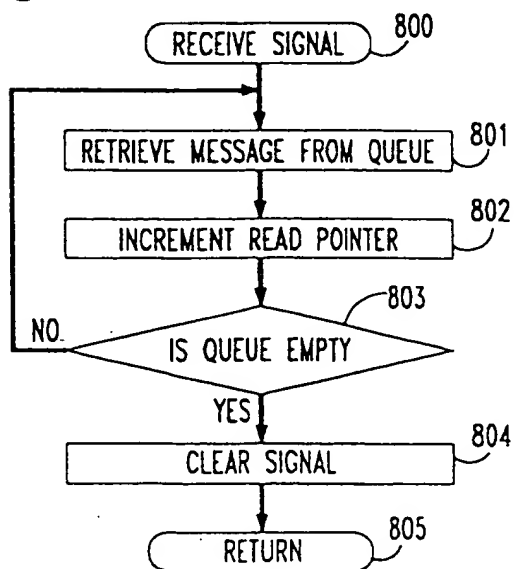


FIG. 8





Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number : **0 592 117 A3**

(12)

EUROPEAN PATENT APPLICATION

(21) Application number : 93307297.7

(51) Int. Cl.⁵ : **G06F 9/46**

(22) Date of filing : 16.09.93

(30) Priority : 24.09.92 US 951488

(43) Date of publication of application :
13.04.94 Bulletin 94/15

(84) Designated Contracting States :
DE ES FR GB IT

(88) Date of deferred publication of search report :
20.04.94 Bulletin 94/16

(71) Applicant : **AMERICAN TELEPHONE AND
TELEGRAPH COMPANY**
32 Avenue of the Americas
New York, NY 10013-2412 (US)


(72) Inventor : **Allen, James Riley**
1364 Deep Run Road
Naperville, Illinois 60540 (US)

(74) Representative : **Buckley, Christopher Simon
Thirsk et al.**
**AT & T (UK) LTD., AT & T Intellectual Property
Division, 5 Morningside Road
Woodford Green, Essex IG8 0TU (GB)**

(54) Asynchronous inter-process communications arrangement

(57) To overcome the lack of support in time-sharing and uniprocessor operating systems such as the UNIX® operating system for real-time and multiprocessor applications, there is provided an asynchronous inter-process communications capability that can be grafted onto the operating systems. Communicating processes (100, 103) communicate via datagram messages through logical asynchronous inter-process communications links (110) each having a synchronous segment (101) and an asynchronous segment (102). The links include a message-serving hub process (11) that communicates in a synchronous, buffer and semaphore based, manner with processes (100) that are message senders, and communicates in an asynchronous, queue and signals based, manner with processes (103) that are message destinations. The hub process may be implemented at any process level of the operating system.

FIG. 2



THIS PAGE BLANK (USPTO)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 93 30 7297

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.5)
X	COMPUTER TECHNOLOGY REVIEW vol. 11, no. 6, May 1991, LOS ANGELES US pages 23 - 27 XP230823 PETER TAIT 'Message passing holds the key to distributed computing' * the whole document *	1,2,4,8, 10	G06F9/46
Y	---	3,6,9 5,7	
A	---	3,6,9	
Y	COMPUTER vol. 22, no. 7, July 1989, LONG BEACH US pages 66 - 77 XP39035 ANNE DISHING 'A survey of synchronization methods for parallel computers' * page 68, middle column, line 1 - line 23 * * page 71, middle column, line 8 - right column, line 7 * * page 72, right column, line 57 - page 73, left column, line 6 * * page 73, right column, line 2 - line 15 *	3,6,9	
A	---	5,7	TECHNICAL FIELDS SEARCHED (Int.Cl.5) G06F
A	COMMUNICATIONS OF THE ASSOCIATION FOR COMPUTING MACHINERY. vol. 14, no. 10, October 1971, NEW YORK US pages 667 - 668 P.J. COURTOIS ET AL. 'Concurrent control with "readers" and "writers" * abstract; figure 2 * -----	5,7	
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 23 February 1994	Examiner Kingma, Y
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons A : member of the same patent family, corresponding document	

EPO FORM 1503 (01.92) (P04C01)

THIS PAGE BLANK (USPTO)